



**What Happens When Employers are Free to Discriminate?
Evidence from the English Barclays Premier Fantasy Football League**

Alex Bryson

National Institute of Economic and Social Research
Centre for Economic Performance, LSE

Arnaud Chevalier

Royal Holloway, University of London
Centre for the Economics of Education, LSE
Deutsche Sporthochschule (Koeln) and IZA

NIESR DISCUSSION PAPER NUMBER 427

ABSTRACT

Research on employers' hiring discrimination is limited by the unlawfulness of such activity. Consequently, researchers have focused on the intention to hire. Instead, we rely on a virtual labour market, the Fantasy Football Premier League, where employers can freely exercise their taste for racial discrimination in terms of hiring and firing. The setting allows us to eliminate co-worker, consumer-based and statistical discrimination as potential sources of discrimination, thus isolating the effects of taste-based discrimination. We find no evidence of racial discrimination, either in initial hiring or through the season, in a context where employers are fully aware of current and prospective workers' productivity.

Key words: race; labour market discrimination; football

JEL: J15; J23; J24; J71; M51

Acknowledgements

We thank participants at the Work, Pensions and Economics Group (WPEG) Conference 2012 in Sheffield, the European Society of Population Economics 2013 in Aarhus, the Royal Economic Society 2014 in Manchester and seminars at the Centre for Economic Performance and the Deutsche Sporthochschule (Koeln) for valuable comments. Alex Bryson would like to thank the Norwegian Research Council (grant number 202647).

1. INTRODUCTION

In most advanced economies it is unlawful for employers to discriminate against individuals either directly or indirectly on grounds of race, either in their recruitment practices or once an individual has become an employee.¹ Despite anti-discrimination legislation numerous studies suggest that employers do discriminate on grounds of race. One set of studies identifies discrimination by seeking to control for differences across workers in wage regressions, treating remaining racial differences as evidence of discrimination. Another set use audit studies/field experiments to identify employer discrimination in recruitment by submitting made-up curricula vitae differing only by race to real job vacancies. Those studies identify a reduced likelihood of an employer "call back" for black applicants relative to identical candidates whose race can be inferred as "white". However, the first type of study is often limited by the lack of individual level productivity measures, while the second type only identify possible discrimination at one early stage in the hiring process.

There are various sources of discrimination in the labour market. Becker (1957) emphasises employers' taste-based discrimination, but also notes that customers and employees may discriminate. Arrow (1972, 1973) and Phelps (1972) emphasise asymmetry of information as a source of discrimination.² In the absence of information on the productivity of prospective employees, employers rely on group statistics as evidence on which to base their hiring decisions. As noted by Levitt (2004: 433) "in

¹ There are exceptions. Policies of "positive discrimination" and "affirmative action" may permit employers to discriminate in favour of minority groups who are identified as particularly disadvantaged in the labour market.

² For a recent review of these theories and their ability to account for racial differences in labour market outcomes see Lang and Lehmann (2011). As the authors point out, the presence of prejudiced employers can lead to differential impact of search frictions across race groups, helping to explain black-white differences in equilibrium unemployment and employment.

general, empirical tests have a difficult time distinguishing between taste-based and information-based models of discrimination”.

This paper contributes to the literature in this dimension as we are able to differentiate between taste-based and statistical discrimination theories in a labour market context. Charles and Guryan (2008) provide evidence that at the State level, prejudice matters and negatively affects the wages of black workers. In our set-up we can directly identify employer taste-based discrimination since productivity is measured at no cost for each potential employee, wages are set outside the firm, firm production is simply additive in each worker and firms do not have customers. As such, all the potential channels of discrimination are blocked apart from employer taste-based discrimination, which in this environment is not unlawful (information regarding the workers employed remains private to the firm) and can be fully satisfied. The environment is similar to a laboratory experiment but with the advantage that employers’ decisions are observed for months. We observe the dynamic of discriminatory behaviour, as more information on the potential workers’ productivity is revealed overtime, and workers can be sacked/hired at low cost.

Our study is based on the analysis of the virtual labour market created by the Fantasy Premier League (FPL), an on-line competition based on the English football Premier League entered by about 2.5 million individuals each season.³ The aim of participants in the FPL is to assemble and manage the best performing team in the fantasy league, something that is achieved by accumulating points related to the performance of professional footballers on the pitch. Employers hire their initial team and can

³ For further details check: <http://fantasy.premierleague.com/>

subsequently buy and sell players in any game week during the 38 week season. This initial squad choice at the start of the season is the focus of our first set of analyses. Our first dependent variable is the number of times a player is hired. For subsequent weeks the data at our disposal are the weekly net demand for a given player after productivity information and price have been updated.

Our analyses examine what role, if any, workers' race plays in employers' choice of their squad at the start of the season and their subsequent decisions to recruit and retain players each week of the season, conditional on their productivity and other factors. While the popular press has reported cases of racial discrimination between footballers and from fans in English football, we find no racial differential in employers' choice of their squad either for the start of the season or in subsequent weeks as more accurate productivity information becomes available. These results suggest taste-based discrimination is absent when fantasy league employers hire workers.

The remainder of the paper is set out as follows. Section Two reviews the previous literature on racial discrimination in the labour market and other markets. Section Three presents our data and the institutional setting for the empirical analysis. Section Four outlines our empirical strategy. Section Five presents results and Section Six concludes.

2. PREVIOUS LITERATURE

There are a variety of reasons as to why employers discriminate on grounds of race. Profit maximising employers may exploit the labour market vulnerability of certain

groups of workers, such as illegal migrants, by offering them employment at lower wages than other workers vying for the same positions, leading to an increased propensity to be taken on but at a lower wage than might otherwise be the case. Alternatively, employers may have a "taste" for employing individuals "like" themselves, in which case white employers are engaged in what Becker (1957) termed "taste-based discrimination" when they offer jobs to whites before non-whites regardless of their aptitude for the job. In such circumstances, employers may pay a price for their taste-based discrimination if their recruitment or promotion procedures are based on skin colour rather than aptitude or productivity. Becker argued that the price of such discrimination is not sustainable in the long-run in competitive markets. Another possibility is that, in the absence of information on prospective employees' worth, employers may judge the quality of applicants based on group characteristics, such as age or race, resulting in what has been termed "statistical discrimination" (Arrow, 1972, 1973; Phelps, 1972).

Studies capturing perceptions of racial discrimination in the labour market suggest it remains commonplace, a finding which is supported by depth interviews with employers themselves (see Pager and Shepherd (2008) for a review). Reviewing audit studies which identify racial differentials in hiring rates, Altonji and Blank (1999: 3194) conclude: "the studies to date generally suggest that hiring discrimination continues to occur". These studies typically send CVs that differ only in the implied race of the applicant to real job openings. Bertrand and Mullainathan (2004) for example, find substantial racial discrimination in call-backs, which is uniform across occupation, industry and employer size. This finding has been replicated in studies using a similar

set-up.⁴ These field experiment studies are clean in the sense that they are able to isolate the role of race on hiring through the manipulation of curricula vitae, but they suffer from the fact that no actual hiring takes place. What they observe instead are 'call-backs' or offers; i.e. discrimination at an early stage of the recruitment process. In our data real hires occur.⁵ In this sense, our setting is similar to Goldin and Rouse's (2000) study which used real auditions for musicians and found women were more likely to be hired in a blind audition when the employers could not see the sex of the musician. In a laboratory experiment Dovidio and Gaertner (2000) find employers discriminate on racial grounds, but only in the case of applicants whose qualifications mean the hiring decision is a difficult one.⁶

It is only recently that discrimination studies have sought to distinguish between statistical and taste-based discrimination. List (2004) shows sports-card traders from minority groups receive lower initial and final offers than those from majority groups. In four complementary follow-up experiments exogenously manipulating information on race he finds the observable differences in treatment are due to statistical discrimination. Similarly, Zussman (2013) experimentally manipulates information on the race of online car buyers and sellers in Israel and finds discrimination against Arab buyers and sellers is statistical rather than taste-based. Doleac and Stein (2013) adopt similar techniques in the online sale of iPods in the United States: they conclude that discrimination against black sellers is due to statistical discrimination rather than taste-based discrimination.

⁴ See Bendick (2007) for a review of audit studies providing evidence of employment discrimination on grounds of race.

⁵ As in the case of most audit and correspondence studies, we lack information on the identity of the employer so we are unable to test for exophobia and endopheilia (Feld et al., 2103).

⁶ Charness and Kuhn (2011) summarize the experimental literature on discrimination. In the laboratory, racial discrimination appears to stem mostly from stereotyping.

Both Zussman (2013) and Doleac and Stein (2013) point to distrust across groups as the cause of such behaviour, which is also consistent with Pope and Sydnor's (2011) evidence that black borrowers face higher interest rates in peer to peer lending. To our knowledge, our study is the first to seek to isolate taste-based discrimination in a labour market setting.

Turning to the sports literature on racial discrimination, the consensus is that racial discrimination has declined over time. Reviewing the wage discrimination literature for the United States, Rosen and Sanderson (2001: F58) suggest that the discrimination which "was easily detected in the initial studies of the 1960s and 1970s...had mostly disappeared by the 1990s...It is difficult to find a negative coefficient on race in US data these days". Kahn (2009) suggests that the racial discrimination on compensation in basketball found in early studies disappeared over time, although there is some evidence of an unexplained black-white salary shortfall among elite players (Hamilton, 1997). Further, two studies on hiring decisions for marginal workers suggest no racial bias against players or coaches in basketball (Brown et al., 1991; Kahn, 2006). However, in their review of the sports literature through to the late 1990s Altonji and Blank (1999: 3196) argue that there is evidence of salary discrimination, especially in professional basketball, some customer discrimination against minority players, and "some hiring discrimination, although these results depend on the sport and position [of the player on the field]".

There also appears to be some diminution in the degree of customer discrimination. An early study identified racial discrimination in the value of baseball cards traded by individual collectors (Nardinelli and Simon, 1990). The price paid for

black and Hispanic retired baseball players is lower than that for whites conditional on career performance statistics.⁷ However, a more recent paper finds no such price differential (Bodvarsson and Brastow, 1999).

Racial discrimination may be less apparent than it used to be because black players have been integrated into North American professional sports. Goff et al. (2002) treat the integration of black players into North American baseball and basketball as akin to the diffusion of a productivity-enhancing technology. Consistent with this proposition they show black players were more productive than white players during the quarter century over which sports moved from a segregated to an integrated equilibrium. The productivity differential dissipates post-diffusion.

Most of the empirical studies of racial discrimination focus on North American labour markets, especially the sports literature. However, there is one study that focuses directly on racial discrimination in English professional soccer. Szymanski (2000) shows that teams with a higher share of black players have higher performance controlling for payroll expenditures, a finding which is consistent with racial discrimination. Whereas Szymanski uses payroll expenditures to proxy for talent, we have direct match-by-match time-varying data on individuals' on-field labour productivity, measured across a variety of dimensions.

⁷ Unlike their study which captures customer discrimination, our private individuals are picking players in order to win and they are in competition with others.

3. DATA AND INSTITUTIONAL SETTING

3.1 Institutional Set-up

We analyse the virtual market of the Fantasy Premier League (FPL); an on-line game based on the Barclays Premier League, which is the top flight of professional football in England. FPL is played by about 2.5 million individuals who sign up to play the game in the course of a season. Participation in the FPL is free. On subscribing, participants are given a fictional budget of £100 million from which they must purchase a squad of fifteen professional footballers playing in the league⁸. The price of players is set by the FPL and reviewed every week. Like in a real firm, different positions must be filled. Here, a team must consist of two goalkeepers, five defenders, five midfielders and three strikers. These players are real footballers playing professional football in the Barclays Premier League. Participants in the FPL are employers in the sense that they buy and sell the players they need in order to produce points and win the league. The overall winner is the team with the most points at the end of the season, or over a month for the monthly prizes. It is also possible to enter teams in private leagues, so as to compete amongst friends. These competitions create incentives for FPL participants to maximize the number of points scored throughout the season, even when an overall win is no longer possible.

As well as selecting their initial squad, employers are able to fire and hire new workers after each game, subject to budget constraints. The cost of a hire is the value of the incoming player plus the gap between the value of the outgoing player on the open

⁸ There are roughly 600 players to choose from in a given season. They cannot hire more than three footballers playing for the same club in the Premier league.

market and the value the employer recovers on sale (which is not the full market price).⁹ Employers are permitted one transfer per week which does not affect their accumulated points total. Any additional transfers entail a deduction of four points, which must be added to the financial cost of making a transfer. Once a year, FPL participants are allowed to hire an unlimited number of footballers with no points penalty.

The points scoring system, i.e. the production function of our firms, is presented in Appendix Table A1. Footballers score points for playing in that particular week, for the time spent on the pitch, and for the actions they perform (positive points for goal scoring, assists and the like, and negative points for own goals, and disciplinary offences leading to red and yellow cards) and bonus points for overall performance. Bonus points are awarded to the best three players in each game, again based on some pre-determined metrics (see Table A1). Productivity is thus objectively measured.

Demand for particular players reflects what is known about their on-field performance, i.e. their productivity, as well as their cost to the employer – as determined by the market value of the player set by FPL – and employers’ personal preferences. Employers have excellent information on each player’s on-field performance across the dimensions described in Table A1, both in previous years when the season starts and in the past games as the season progresses. As such, information about the productivity of each potential worker improves over time, and is available to all firms. Importantly to study racial discrimination, the summary information presented for each footballer contains a colour picture, as well as his position, team, the proportion of other employers

⁹ There is a gap between the buying and selling prices of players. This margin is half of the difference between the current price and the price at which the player was bought; this can be thought as a tax on the value added. As such, transferring players has a financial cost and leads to a reduction in the firms’ budget. So firms may not always optimize their teams and may refrain from using their weekly transfer.

who have that player in their squad, his performance in recent matches, his current market value, upcoming fixtures, the number of games the player is scheduled to play this week (usually one, but sometimes, zero or two) and an injuries update. This is illustrated in Figure A1 for Emmanuel Adebayor. These pictures were used to determine the race of each footballer, using a dichotomous categorisation: white, non-white.¹⁰

Seven features of this setting mean that we can recover more precise estimates of racial discrimination in relation to hiring than is possible in other settings. First, we are able to identify the effects of taste-based employer discrimination, as opposed to the effects of customer, co-worker or statistical discrimination. There is no possibility of customer discrimination since employers do not have clients.¹¹ Team production is simply additive in individual workers' production; the absence of co-worker relations means there is no co-worker discrimination. It also means that we can ignore the importance of productivity spill-over across footballers, which would complicate recruitment and retention decisions by bringing in factors other than individual talent. Employers have access to very comprehensive weekly data on the productivity histories of all workers in the industry, together with their market prices, so their information set regarding worker value is much richer than would ordinarily be the case. Since the productivity of each worker (not only of employees) is perfectly known, at least as the season progresses, there is no scope for statistical discrimination. Employers also know the skin colour of all workers in the population of potential recruits: it is not just proxied by name as in most field experiments (Bertrand et al, Mulainathan, 2004, for example).

¹⁰ There are few players defined as other races – for the analysis they have been grouped with black players. This ethnicity variable was defined by three individuals and race determined by a simple majority rule.

¹¹ In a real world setting customer discrimination can affect team selection via crowd attendance at games (Bryson et al., 2014).

They are therefore able to account fully for the performance of workers and their race in decisions concerning recruitment and retention. Thus, if there are any indications of racial bias, they are unlikely to reflect anything but employer taste-based discrimination.

Second, employers are free to discriminate in their hiring and firing behaviour. In this sense we are "turning back the clock" to a time when employers faced no legal impediments to discrimination. Therefore the costs of discrimination are low and we can thus identify an unbiased taste for discrimination.

Third, the setting is a single occupation in a single industry, so workers are perfectly substitutable for one another (within a position on the field), and the full productivity history of workers is available at no cost; i.e. there is no monitoring cost. Thus, this study can overcome the problem that, in many observational studies, it is difficult to compare "like-for-like" workers.

Fourth, as in a laboratory experiment, the firms are identical in size, technology and initial budget. In assembling the workers required by the firm, employers must fill identical job slots within the firm. At the beginning of the year all employers face the same budget constraint, so their ability to recruit a mix of more and less talented workers is identical at the outset, although budgets vary as the season progresses due to value added (destroyed) when selling workers.

Fifth, although firms are in competition with one another, workers are able to work at more than one firm simultaneously so that firms are not in direct competition with one another for worker talent. Thus, in principle, all workers are available for hire, subject to firms' budget constraints.

Sixth, workers have no say in the firms they join and can only exit if fired, so there is no selection of workers into more (less) discriminating firms.¹²

Finally, employers are price takers: the price of recruiting individual workers, i.e. a sign-on fee, varies substantially but individual employers are unable to influence these prices. Prices attached to workers are exogenous to the firm, but relate very strongly to worker performance, as we shall see. As such, firms cannot exploit minority workers by offering them lower sign on fees. Once signed, there are no wages in our set up.

3.2 *Does the FPL function like a normal labour market?*

Before we investigate the racial differences in player hires we need to establish whether the FPL functions like a labour market. Evidence to this effect is presented in Figure 1. Panel A shows that better performing players in the previous season are valued more highly at the start of the new season. Panel B shows that demand for players, as measured by the number of times a player is picked for initial squads, rises steeply with performance in the previous season. The steep non-monotonic rise in price and demand for high-performing players evident in Panels A and B is consistent with a market for superstars, as originally conceived by Rosen (1981). The discontinuity is at 180 points, which represents the top 10% performers. Unsurprisingly, the initial demand/initial price plot follows a similar, albeit smoother, trajectory (Panel C). The remaining two panels (D and E) look at the same relationships as the season progresses. Previous week productivity is positively correlated with demand and change in price. Overall, the FPL

¹² Worker selection based on perceptions of discriminatory tendencies in particular occupations or among certain employers may contribute to wage discrimination. For example, Plug et al. (2014) find gays and lesbians in Australia shy away from more prejudiced occupations.

appears to behave like a labour market, where more productive workers are in higher demand and command higher fees

[FIGURE 1]

3.3: Data Description

Each week, the FPL participants select the 11 players from their 15 man squad who will score points for their fantasy team depending upon their performance in real football games played that week, as well as a captain whose productivity will be doubled. The data available to us only cover the number of employers who have a given footballer on their books, not whether this player has been selected to score points for the team. In that sense, our set up is similar to audit studies in that we have information on employees not employers. We share with these studies the lack of knowledge on the characteristics of employers, but survey evidence amongst the population of fans attending Premier League games during the season 2008/09 suggest that 8% were non-white, while ethnic minorities represent 14% of the UK population.¹³ Football fans in England are thus disproportionately white. However, while based on the English Premier League, the FPL is open to anybody.

We have match-by-match data on players' performance as indicated by the FPL scoring system of all footballers registered in one of the 24 football teams playing in the Premier League in the three seasons 2008/09, 2009/10 and 2010/11.¹⁴ This gives us a

¹³ http://www.epfl-europeanleagues.com/changing_face.htm and <http://www.ons.gov.uk/ons/rel/census/2011-census/key-statistics-for-local-authorities-in-england-and-wales/rpt-ethnicity.html>.

¹⁴ The English Premier League is composed of 20 clubs each season. At the end of the season, the bottom three are dropped (relegated) and replaced by the best performing teams from the Championship league. We only observe 24 teams since Newcastle United and West Bromwich Albion were relegated at the end of the 2008/09 season but gained promotion to play in the 2010/11 season.

total of 1,767 footballers, about 600 per season. Since the Premier league starts before the transfer window for professional footballers is closed, not all footballers are available to be picked for the first game. Subsequently, professional teams can transfer players in January, and players from their academy may join the league during the season. As such, only 1,327 players are available for pick in week 1.¹⁵ Each club plays 38 games in the season. We have thus an unbalanced panel of 60,086 player-match observations. Thirty-six per cent of potential employees are non-white, two and half times above the percentage of non-white in the English population as a whole.¹⁶

The player performance data available to all employers comes from FPL, which runs the fantasy league. Players' productivity is based on rudimentary objectively verifiable data of their performance in a game, as explained above and in Appendix Table A1. Individual points in a given week range between -6 and +32 with a mean of 1.50.¹⁷ Points scored by players is of interest in its own right since we can establish to what extent there are any racial differences in the productivity of players. Table 1 and Figure 2D report no difference in weekly performance by race. Over the full season, total productivity is also very similar between white and non-white players, with the average player scoring about 40 points. Note however that the superstar footballers are disproportionately white: whites represent 78% of the top 1% performers (Figure 2A).

[TABLE 1]

[FIGURE 2]

¹⁵ That is 444 players in season 2008/09 and 2009/10 and 239 in season 2010/11.

¹⁶ Szymanski (2000: 597) notes that there were only 4 black players playing in the 38 English professional football clubs in his data in 1974. By 1993 black players were much more common, accounting for around 8 per cent of his sample.

¹⁷ Not all footballers play in a given week, so the mode score is actually zero.

Time on the pitch enters our performance metric. But this could itself be a function, in part, of racial discrimination among "real world" Premiership coaches if their decisions regarding who to play and how long to play them for are racially biased, either because they are responding to customer preferences for white players, or because they are indulging their own taste-based discrimination or statistical discrimination. There is, however, little evidence of racial differences in playing times; whites are marginally less likely to play but conditional on playing, play for a few more minutes. These two inputs appear to compensate for one other so they are unlikely to contribute to any variation in overall productivity.

We find little evidence of differences in productivity between players of different races, supporting Goff et al.'s (2002) observation that top-flight professional sports are racially integrated. The second main determinant of demand for a player is price. Prices are set by the FPL and appear to be a function of performance and net demand (see below). At the beginning of the season, price ranges between £3.9 million and £14 million (Christiano Ronaldo) and, on average, non-white players are marginally more expensive at the outset. Indeed the distribution of initial price is slightly shifted to the right for non-white players (Figure 2B) This is surprising considering that there is no difference in productivity but may reflect the fact that players differ along racial line on observable characteristics. Non-whites are more likely to have been present in the league last season, as such there is less uncertainty about their productivity. They are also more likely to be non-British nationals, more likely to have played for their national team and to play at least once in the season. More importantly, they play in different positions. Non-whites are more likely to play as forwards, the most expensive players on average

and they are least likely to be goalkeepers, the least expensive players on average.¹⁸ Below we come back to this issue of whether the set prices are fair. After the initial valuation, players' prices are reviewed on a weekly basis. In any week, 80% of prices remain the same, and weekly price adjustments are in general small. These distributions are almost identical by race (Figure 2E) but non-whites lose marginally more value than white players in any given week (-0.007 vs -0.005).

Table 1 also reports the mean of the two outcomes of interest: initial demand and weekly net flow. The number of times a player is chosen by employers in their initial squad is only available for the 2009 and 2010 seasons. On average, there are 70,000 picks for a white player and 60,000 for a non-white but this difference is not statistically significant. The distributions of initial demand are almost indistinguishable but note that at the top end, there are only white players. There is no difference in net transfers and the distributions are perfectly super-imposed.¹⁹ Based on these descriptive statistics there appears to be little evidence of discriminatory behaviour amongst FPL participants.

4. EMPIRICAL STRATEGY

We investigate two potential dimensions along which racial discrimination may occur, namely initial hires at the start of the season and transfers during the course of the season. There may be discrimination in the labour market if, conditional on performance, employers are simply less likely to purchase non-white players at a given price. Since all

¹⁸ Lang and Lehmann (2011) discuss constraints on employer ability to discriminate when potential working populations are heavily skewed racially. Even for the most racially biased position, goalkeeper, non-whites represent 10% of the potential employees.

¹⁹ The mean net transfer is positive as some participants join the FPL during the season. Those late participants buy 15 players and do not sell any, leading to a small positive net demand overall.

other channels of discrimination are blocked, racial differences in hires will arise through taste-based discrimination in the way Becker envisaged (see Altonji and Blank, 1999: 3170). As Kahn (2009: 14-15) notes, identification of racial bias in hiring and firing decisions is best investigated using performance differences of marginal workers, as opposed to the average worker because only the former is informative about the margin where the hiring/firing decision is made. This is precisely what we observe in our data since all players are available for hire by all employers at any point in time, and can be dismissed with the low dismissal costs described in Section Three.

First we investigate whether there are any racial differences in employers' squad choice for the start of the season and, if so, whether they can be accounted for by performance in the previous season, initial price and other observable characteristic of the player. If a player is new to the league and has no information on performance in the prior season, values for previous season metrics are set to zero and the observations identified with a dummy variable set to unity for players new to the league. This category is a mix of young players, players who used to play in lower leagues (especially those of promoted teams) and players who previously played abroad. We thus estimate a regression of the (log) number of initial picks for each professional footballer, accounting for race, a quadratic in price, a quadratic in previous year performance and the player's characteristics.

$$\ln Pick_i = \beta_1 Race_i + \beta_2 Fees_i + \beta_3 Fees_i^2 + \beta_4 Perf_i + \beta_5 Perf_i^2 + \beta X_i + \varepsilon_i \quad (1)$$

Standard errors are clustered at the individual level, so as to capture correlation in the popularity of a player between seasons. Additional robustness checks estimate this model using quantile regressions rather than ordinary least square.

Second, we investigate the determinants of employers' net demand for players, that is, the difference between the number of hires and the number of fires a player is subject to each week, as participants to the FPL are allowed to revise their team.

There are three reasons to suspect that discriminatory behaviour of employers may diminish over the course of the season. First, as Antonovic et al. (2005: 923) note in the context of *The Weakest Link* TV show, "the implicit cost of taste-based discrimination rises as the game progresses (because one's probability of winning the game is higher in later rounds) discriminatory outcomes due to taste-based discrimination should diminish over time". Intensified competition should limit such behaviour in the FPL for the same reasons. Second, whereas at the start of the season employers must rely on player performance information from the previous season, employer information about player performance is continually up-dated throughout the season such that, if there is any uncertainty regarding productivity at the outset that could induce some statistical discrimination, this dissipates over time as employers observe player's "form", including that of footballers they have not selected such that the information on productivity is perfect and covers the full set of employees and potential employees. Third, there may be attrition in the employers remaining active in the league if, for example, it is those who are most committed to winning, or have the greatest chance of winning, who continue to hire and fire to the end of the season. Figure 3 shows movement in the net demand for players over the course of the 38 week season for the pooled seasons. It is apparent that market activity declines over the course of the season, perhaps reflecting falling effort as most employers find they are unable to win, or growing employer

perceptions that they have optimised in the face of budget constraints.²⁰ The sudden drop before week 20 is likely due to the Christmas holiday season. If we assume that discriminatory behaviour is costlier for these employers still playing towards the end of the season (since under-performance via discrimination is most keenly felt by these employers) differential attrition should mean discrimination declines over the season. For all these reasons we should expect diminishing discriminatory behaviour as the season progresses. As well as flexibly controlling for time, we run robustness check splitting the season into four periods.

[FIGURE 3]

We isolate race-related differences with a dummy variable identifying non-white players. The specification includes the same individual characteristics as in (1), plus indicators of weeks of play, to capture potential fluctuation in participants' interest through the season, and an indicator on the number of games scheduled to play during the week (usually one, but sometimes zero or two). Fees are the updated values after last performance. The measures of productivity capture various time dimensions. A priori, employers may use different productivity metrics, and face a trade-off between using recent productivity which may capture some random luck component, or use a long term average which may fail to capture recent change in player's form. We include three, covering different periods: the immediate productivity (last period performance), the three previous periods, and the overall performance since the season started. Considering

²⁰ We are unable to identify individual employers in our data so we are unable to assess the direct effects of employer attrition.

that we have repeat observation for players, standard errors are again clustered at the player level.²¹

$$\Delta Pick_{it} = \beta_1 Race_i + \beta_2 Fees_{it} + \beta_3 Perf_{it-1} + \beta_4 Perf_{it-2/4} + \beta_5 Perf_{i1/t-1} + \beta X_{it} + f(t) + \varepsilon_{it} \quad (2)$$

Finally, to absorb any remaining characteristics of the player that may not be captured by these observables, we also estimate a player/season level fixed effect. The race effect is then identified through its interaction with performance and price; i.e. are non-white players more likely to be traded than white players at a given level of recent performance or price.

$$\begin{aligned} \Delta Pick_{it} = & \beta_2 Fees_{it} + \beta_3 Perf_{it-1} + \beta_4 Perf_{it-2/4} + \beta_5 Perf_{i1/t-1} + \\ & \gamma_2 Fees_{it} * Race_i + \gamma_3 Perf_{it-1} * Race_i + \gamma_4 Perf_{it-2/4} * Race_i + \gamma_5 Perf_{i1/t-1} * \\ & Race_i + \beta X_{it} + f(t) + \mu_i + \varepsilon_{it} \end{aligned} \quad (3)$$

5. RESULTS

5.1 Graphical Evidence

Figure 4 presents information on the relationship between performance, price and demand for players by race, both at the start of the season and during the season. As shown from Figure 2, it is notable that there are only white players towards the very top of the rankings in terms of last season's performance, raising questions as to whether one can recover a white-nonwhite differential in demand conditional on performance at the very top of the performance scale. Where there appears to be common support the

²¹ We also estimated models with standard errors clustered at the player/season level, but results were very similar and are not reported here.

association between initial demand for players at the start of the season and their performance in the previous season appears very similar by race (Panel A). Note that while the relation between performance and pick is linear there is a definite kink for top performers. The same is also true during the course of the season, with demand rising with points for both whites and nonwhites in a similar fashion, at least up to the higher reaches of the points distribution, after which point the number of player observations is very small (Panel C). Similarly, throughout most of the points distribution the association between player performance and price is similar by race, both at the start of the season (Panel B) and during the course of the season (Panel D), but while the pricing of performance is linear for whites, for non-whites it tails off at very high weekly productivity levels.

[FIGURE 4]

5.2 *Racial Differences in Outcomes Other than Employer Demand for Players*

As Table 1 shows, white and non-white players differ in observable traits in ways which may also influence players' propensity to be chosen for starting squads. It is therefore sensible to move to a multivariate framework in which we can control for racial differences in observable traits which might obscure employers' propensity to discriminate. Table 2 presents results for eight different outcomes - four related to players at the beginning of the season, and four outcomes that unfold during the course of the season. It is apparent that, conditioning on the observable player and team traits displayed at the bottom of the table, there are no racial differentials with respect to the various

dimensions of footballers' productivity.²² The only statistically significant differential is that non-white players are valued at around 2 per cent less than observationally equivalent white players at the beginning of the season (row 1), something that remains over the season, since there is no difference in price change as the season unfolds (row 5).²³ Since we control for initial price in most regressions, this differential will not affect our estimates of discrimination in hiring.

[TABLE 2]

5.3 *Initial Squad Selection*

Table 3 presents our first main set of findings on log hires at the start of the season. There is no racial differential in initial hires when comparing raw hiring rates (column 1), a finding that holds conditioning on player's position, age, birth place, international caps status and club playing for in the Premier League (column 2). However, initial price plays a big part in whether a player is selected by employers for their initial squad, as we saw in Figure 1 Panel C. When we condition on initial price we find non-white players are *more* likely to be selected than white players. This positive discrimination is likely related to the initial under-pricing of non-white players. However, the non-white coefficient falls in size with the addition of player performance in the previous season, such that the racial differential is no longer statistically significant (column 5). When including both price and productivity measures, initial demand is driven by past productivity but not price. However, players for which past productivity is unknown, because they did not play in the Premier League last season, are also more

²² For points accrued this week and minute played conditioning on having played does not alter these conclusions.

²³ We are not privy to the algorithm used by FPL to determine prices.

popular. This may be driven by well-known foreign players joining the league and benefiting from a halo effect.

[TABLE 3]

The distribution of initial pick is characterised by a large bulk of players with very low number of hirings and a long tail of super-stars attracting a large fraction of participants – note that the pricing is such that participants can only afford a few super-stars due to their budget constraint. Thus, we assess whether the racial differential in hiring differs at different points of the hiring distribution. Table 4 reports estimates of quantile regressions estimated for the following quantile: 10, 25, 50, 75 and 90. While the point estimates are consistent with some discrimination happening amongst the superstars, these differences are not statistically significant. Note, that the positive effect of not having a productivity record is driven by very popular players.

[TABLE 4]

Finally, we estimate our favoured specification (Table 3, Column 5) for different subsamples. First, we differentiate for having played in the Premier League previously. Our prior would be that for individuals who haven't played, there is more uncertainty about their productivity such that potential employers may revert to some type of statistical discrimination. Second, we split the sample by previous productivity to test whether variation in the penalty for discrimination faced by employers (in term of lost inputs) affects their discriminatory behaviour. Third, we split the sample by position on the field since productivity and racial composition differ by position. Lastly, we split the panel by nationality, as employers may be more likely to discriminate against non-white foreigners than against non-white nationals. Table 5 displays the results of these tests of

heterogeneity in the racial gap in hiring. In most instances the non-white coefficient is positive but it is only statistically significant in the bottom quartile of last year's performance distribution where non-white players are substantially more likely to be picked. Overall, we find no evidence to suggest that employers are discriminating on grounds of race when making their first hires for the season.

[TABLE 5]

5.4 *Net employer demand for players*

We now turn to net employer demand for players over the course of the season. An advantage of our set-up is that we can observe the dynamics of hiring and firing. This is important since as more information about contemporaneous productivity of employees and prospective employees become available, employers may revise their hiring decisions.

[TABLE 6]

Although net demand for non-white players is lower than it is for white players, the difference is quantitatively small and statistically non-significant in the absence of controls (Table 6, column 1). The non-white coefficient becomes positive but remains statistically non-significant when we introduce controls for player characteristics and three measures of productivity: performance in the last game, performance in the previous three games, and total points accumulated in the games up until that point. Note that hiring/firing decisions appear to be primarily based on the most contemporaneous measure of productivity, even so this measure, to a large extent, may be affected by chance. The coefficient becomes negative when conditioning on player price in column 3

but, again, it is far from statistically significant. The picture remains the same when additionally conditioning on the number of times the player was picked for squads at the start of the season, as this affects the number of future transactions (column 4). In column (5) we test whether employer demand reacts differently to productivity information depending on the race of the player. None of the interactions are significant but the main race coefficient turns statistically significant. However, the effect is still small: being non-white would be related to an decrease of almost 2,000 participants who own a non-white player. This is 0.18 of a standard deviation in net trade. Finally, moving to a fixed effects model (Column 6), a small negative effect between being non-white and accumulated points in the season is apparent when incorporating player fixed effects (column 6) but again this is not economically relevant.

There is little evidence of racial differentials at the mean, and also very little at any of the quantiles in the net transfer distribution (Table 7). Demand for non-white players is statistically significant at the 90th percentile but again the effect is so small as to be economically negligible.

[TABLE 7]

Finally, in Table 8 we test for heterogeneous effects in discrimination across sub-populations using specification 3 of Table 6. First, we restrict the sample to footballers who actually played at least once in the league during the season. Second, we assess whether individuals for whom less information on past productivity was available (because they did not participate in the league in the previous season) are more or less likely to be discriminated against as productivity information is revealed. Third, initial pick is likely to affect the potential subsequent net transfers so we split the sample into

quartiles of initial picks. Fourth, since the racial composition differs by position, we separately estimate the regression for each position. Fifth, we split the sample between national and non-national players. Sixth, we test whether, as expected, discrimination decreases over time for reasons presented above: costs of discrimination, selection of participant, more valid productivity information. On all these dimensions we find little evidence of racial discrimination (Table 8). The only statistically significant effects are the lower likelihood of keeping non-white players among i) those in the third quartile of the distribution of the initial squad demand, and ii) those for which no productivity information was available at the beginning of the season. Both of these effects are small. Although there are theoretical grounds for suspecting the race coefficient to vary over the course of the season, there is no empirical support for this proposition.

[TABLE 8]

6. CONCLUSIONS

We have examined whether there are racial differences in hiring rates in a setting where we can discount the possibility of discrimination on grounds other than taste-based discrimination. This setting is a virtual labour market for professional football players. The players are real: they play each week in England's top professional league, the Barclays Premier League. However, the market for their services is virtual: it is an on-line game where participants operate as employers, buying and selling players in order to win the fantasy football league. The setting is particularly attractive because the football players do not play together once picked, so there is no opportunity for co-worker discrimination; the employers have no customers, so there is no consumer-based

discrimination; and employers have perfect knowledge of individual players' time-varying labour productivity, precluding the possibility of statistical discrimination. What is more, there is no legal bar to employers exercising any taste for discrimination they may possess, making the costs of such discrimination lower than they are in the real world. If taste-based discrimination is an important component in employers' hiring decisions, we would therefore find it here. In addition to our ability to isolate the impact of taste-based discrimination, our setting has a number of other very attractive features including identical firms with identical budgetary constraints; the substitutability of football players for one another (at least within position on the pitch); players are simultaneously available for hire by multiple firms; fees for employees are exogenously given; and players have no choice as to whether they are hired by an employer or not.

We find no evidence of racial discrimination in hires, either at the start of the season, when employers pick their initial squad, nor during the course of the season. Productivity is the main driver of hiring/firing decisions. A rare occurrence of discrimination is in the hiring/firing of employees who are new to the league and for which productivity information at the beginning of the season was not available. However, this effect while statistically significant is very small.

The finding of no discriminatory behavior against non-whites is consistent with other studies which suggest that racial differences in market outcomes are not driven by taste considerations. Instead other factors are at play. For example, a number of studies point to the importance of statistical discrimination (List, 2004; Zussman, 2013; Doleac and Stein, 2013) while others, such as Plug et al. (2014), emphasize the importance of worker sorting based on perceptions of discrimination in certain occupations or among

particular types of employer. Since those channels are closed in our setting, it is perhaps unsurprising that we do not find any evidence of discrimination. Our findings are also consistent with studies which point to a diminution in the extent of racial discrimination in sports on the part of employers and fans. The difference in discriminatory behaviour between the sport context and the general labour market is likely to be driven by the availability of good productivity measures for all possible employees in the sports labour market, something that is clearly not possible in the general labour market.

References

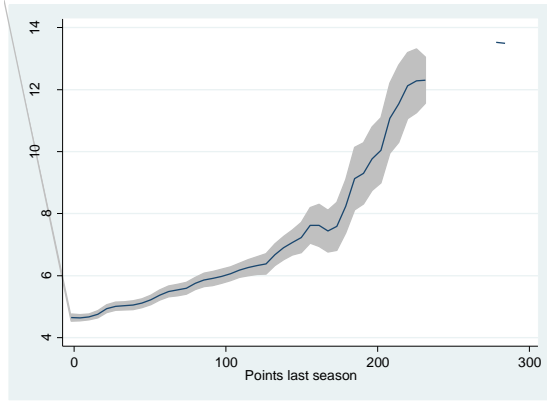
- Altonji, J. G. and Blank, R. M. (1999) "Race and gender in the labor market", Chapter 48, *Handbook of Labor Economics*, vol. 3, 3143-3259
- Antonovics, K., Arcidiacono, P. and Walsh, R. (2005) "Games and Discrimination: Lessons from *The Weakest Link*", *The Journal of Human Resources*, XL, 4: 918-947
- Arrow, K. J. (1972) "Models of job discrimination" in A. H. Pascal (ed.) *Discrimination in Economic Life*, pp. 83-102, Lexington, MA: Lexington Books
- Arrow, K. J. (1973) "The Theory of Discrimination" in O. Ashenfelter and A. Rees (eds.) *Discrimination in Labor Markets*, pp. 3-33, Princeton, NJ, Princeton University Press
- Becker, G. S. (1957) *The Economics of Discrimination*, Chicago: Chicago University Press
- Bendick, M., Jr. (2007) "Situation Testing for Employment Discrimination in the United States of America", *Horizons Strategiques*, 3(5): 17-39
- Bertrand, M. and Mullainathan, S. (2004) "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination", *American Economic Review*, 94(4): 991-1013
- Bodvarsson, O. B. and Brastow, R. T. (1999) "A Test of Employer Discrimination in the NBA", *Contemporary Economic Policy*, 17, 2: 243-244
- Brown, E, Spiro, R. and Keenan, D. (1991) "Wage and Nonwage Discrimination in Professional Basketball: Do Fans Affect It?", *American Journal of Economics and Sociology*, 50, 3: 333-345
- Bryson, A., Rossi, G. and Simmons, R. (2014) "The Migrant Wage Premium in Professional Football: A Superstar Effect?", *Kyklos*, 67, 1: 12-28

- Charles, K. and Guryan, J. (2008) "Prejudice and The Economics of Discrimination", *Journal of Political Economy*, 116(5): 773-809.
- Charness, G. and Kuhn, P. J. (2010) "Lab Labor: What Can Labor Economists Learn from the Lab?", *NBER Working Paper No. 15913*
- Doleac, J. L. and Stein, L. C. D. (2013) "The Visible Hand: Race and Online Market Outcomes", *The Economic Journal*, 123, F469-F492
- Dovidio J. F, and Gaertner, S.L (2000) "Aversive racism and selection decisions", *Psychological Science*, 11(4):315–19
- Feld, J., Salamanca, N. and Hamermesh, D. S. (2013) "Endophilia or Exophobia: Beyond Discrimination", *NBER Working Paper No. 19471*
- Goff, B. L., McCormick, R. E., and Tollison, R. D. (2002) "Racial Integration as an Innovation: Empirical Evidence from Sports Leagues", *The American Economic Review*, 92, 1: 16-26
- Goldin, C. and Rouse, C. (2000) "Orchestrating impartiality: the impact of 'blind' auditions on female musicians", *American Economic Review*, 90, 4: 715-741
- Hamilton, B. H. (1997) "Racial Discrimination and Professional Basketball Salaries in the 1990s", *Applied Economics*, 29, 3: 287-296
- Kahn, L. M. (2006) "Race, Performance, Pay, and Retention Among National Basketball Association Head Coaches", *Journal of Sports Economics*, 7: 119-149
- Kahn, L. M. (2009) "The Economics of Discrimination: Evidence from Baseball", *IZA Discussion Paper No. 3987*
- Lang, K. and Lehmann, J. K. (2011) "Racial Discrimination in the Labor Market: Theory and Empirics", *NBER Working Paper Number 17450*

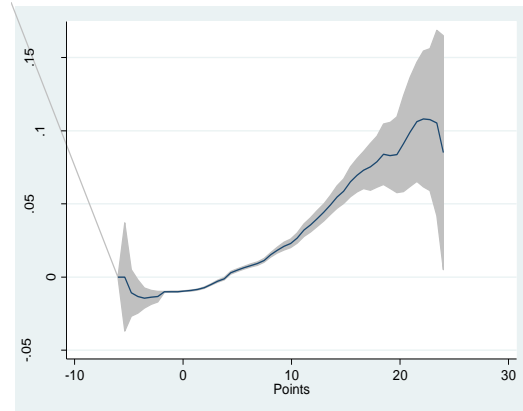
- Levitt, S (2004) "Testing Theories of Discrimination: Evidence from *The Weakest Link*." *Journal of Law and Economics*, 47(2): 431–52
- List, J. A. (2004) "The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field", *The Quarterly Journal of Economics*, 119, 1: 49-89
- Nardinelli, C. and Simon, C. (1990) "Customer racial discrimination in the market for memorabilia: the case of baseball", *The Quarterly Journal of Economics*, 105, 3: 575-595
- Pager, D. and Shepherd, H. (2008) "The Sociology of Discrimination: Racial Discrimination in Employment, Housing, Credit and Consumer Markets", *Annual Review of Sociology*, 1, 34: 181-209
- Plug, E., Webbink, D. and Martin, N. (2014) "Sexual Orientation, Prejudice and Segregation", *Journal of Labor Economics*, 32, 1: 123-159
- Phelps, E. (1972) "The Statistical Theory of Racism and Sexism", *American Economic Review*, 62, 4:659-661
- Pope, D., and Sydnor, J. (2011) "What's in a picture? Evidence of Discrimination from Prosper.com", *Journal of Human Resources*, 46, 1: 53-92
- Rosen, S. (1981) "The economics of superstars", *American Economic Review*. 71; 845–858
- Rosen, S. and Sanderson, A. (2001) "Labour Markets in Professional Sports", *The Economic Journal*, 111, F47-F68
- Szymanski, S. (2000) "A market test of discrimination in the English Soccer Leagues", *Journal of Political Economy*, 108: 590-603
- Zussman, A. (2013) "Ethnic Discrimination: Lessons from the Israeli Online Market for Used Cars", *The Economic Journal*, 123, F433-F468

Figure 1: Does Fantasy League function like a market

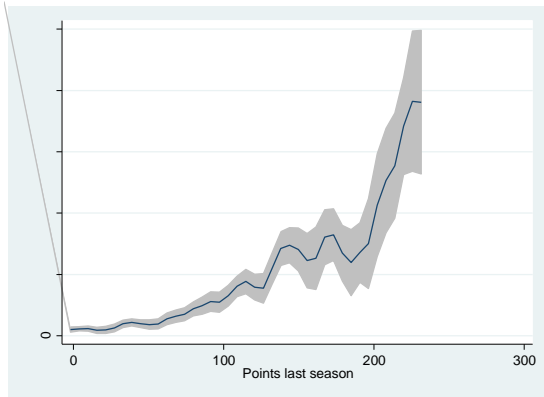
A: Last year Productivity and Initial Price



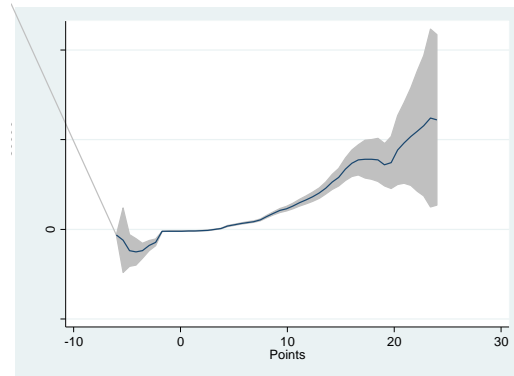
D: Productivity and change in Price:



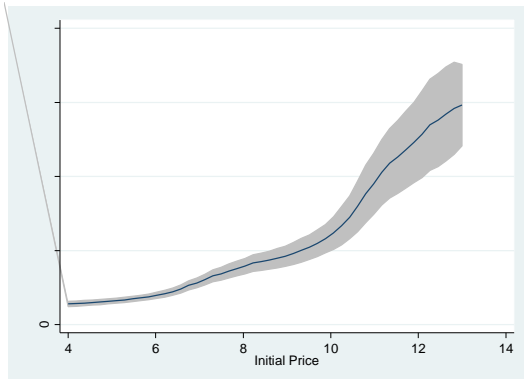
B: Last year productivity and Initial Demand



E: Productivity and net employer demand



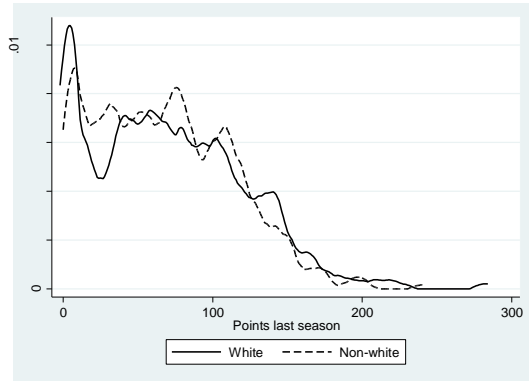
C]: Initial Price and Initial Demand



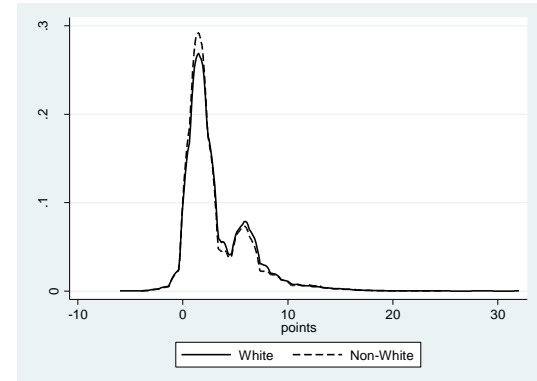
Note: Own calculation based on Fantasy Football League 2008/09, 2009/10 2010/11
 Shaded area represent the 95 per cent confidence interval

Figure 2: Racial Differences in Performance, Price and Demand:

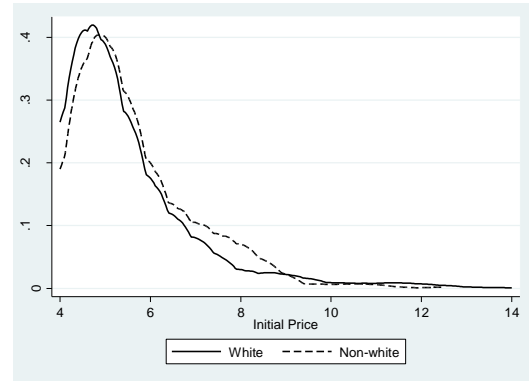
A) Last Year Performance by Race



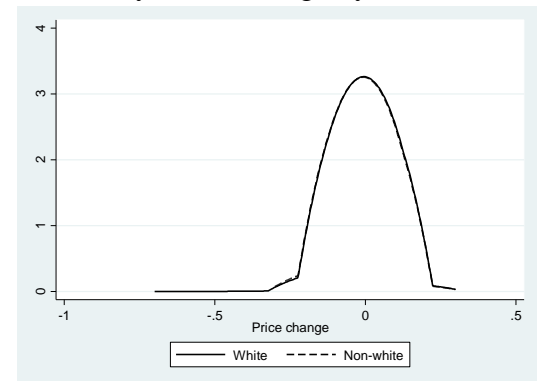
D) Weekly Performance by Race



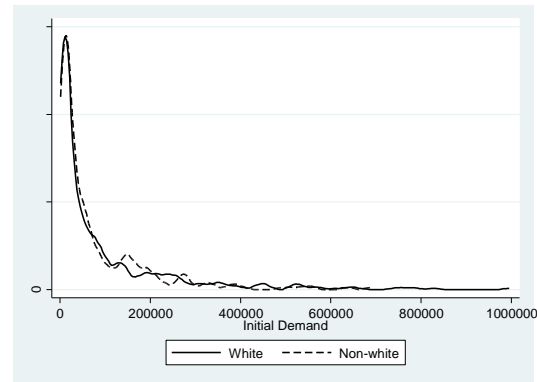
B) Initial Price by Race



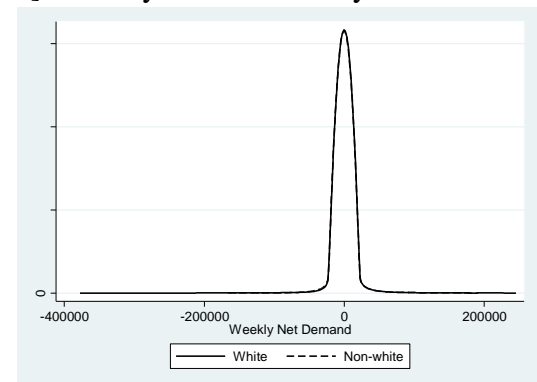
E) Weekly Price Change by Race



C) Initial Demand by Race

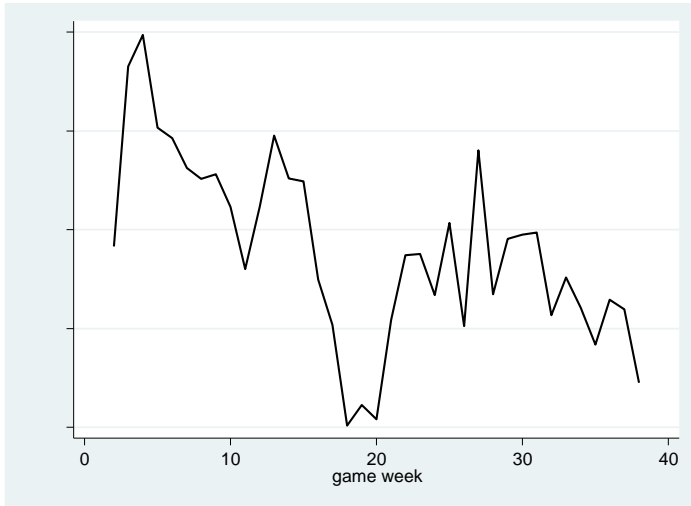


F): Weekly Net Demand by Race



Source: Fantasy Football League 2008/09, 2009/10 2010/11

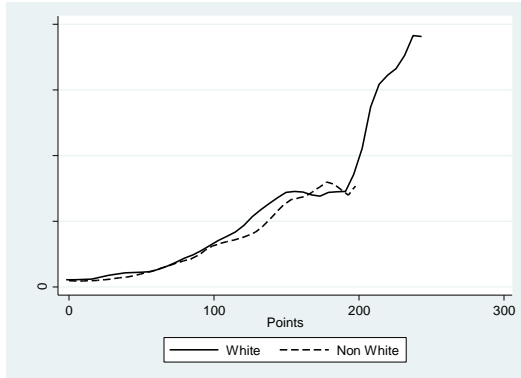
Figure 3: Market activity per week



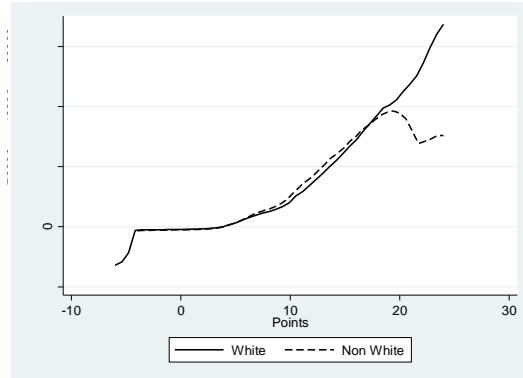
Note: Own calculation based on Fantasy Football League 2008/09, 2009/10 2010/11.
Market activity is the sum of all the absolute net transfers in a given week

Figure 4: Racial Differences in market relationship:

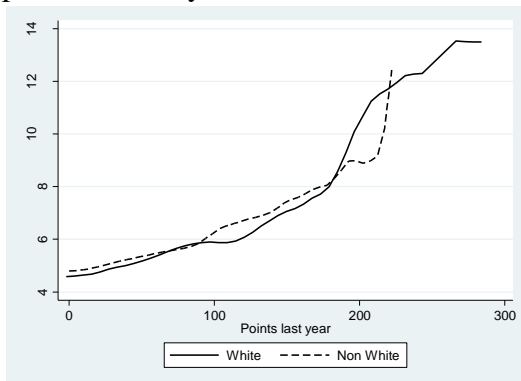
A) First Picks by Race and Previous seasons Points



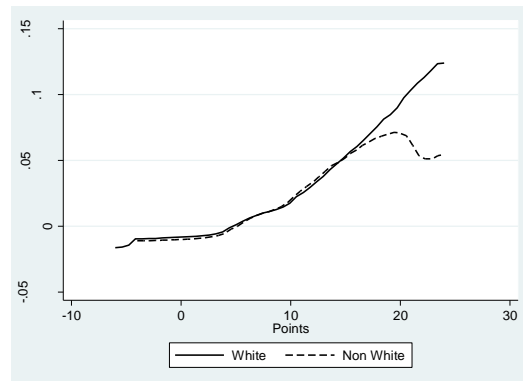
C) Net Demand and Performance



B) Initial Price and last year performance by race



D) Change in Price and Performance



Fantasy Football League 2008/09, 2009/10
2010/11

Table 1: Descriptive Statistics:

Fixed characteristics	White		Non-White	
	mean	s.d.	mean	s.d.
Price at t=1	5.239	1.404	5.403**	1.274
Nbr of picks at t=1 ^a	70779	129493	60226	100267
Points last year	40.751	51.833	42.944	48.517
Top decile points last year ^b	0.104	0.306	0.070**	0.255
Not in league last year	0.367	0.482	0.308**	0.462
Non UK national	0.482	0.500	0.664***	0.473
Age	26.610	5.131	26.079**	4.505
International cap	0.586	0.493	0.651***	0.477
Defender	0.317	0.466	0.333	0.472
Forward	0.147	0.354	0.266***	0.442
Goalkeeper	0.159	0.366	0.030***	0.170
Mid fielder	0.377	0.485	0.371	0.483
Never played in Season	0.160	0.367	0.109***	0.312
Nbr Observations	1125		642	
Weekly varying variable				
Played in Game	0.478	0.500	0.502***	0.500
Minutes played \if played	77.15	0.233	73.54***	0.304
Points	1.508	2.569	1.499	2.525
Change in Price	-0.005	0.037	-0.007***	0.038
Net Transfer	123	11,441	72	10,777
Nbr Observations	38,059		22,027	

Source: Fantasy Football league 2008/2011

***, **, * reflect statistical difference of the means for the two groups at the 99%, 95% and 90% confidence

^a Number of picks is only available for season 2009/10 and 2010/11, the sample size is thus White (751) non White (416).

^b Conditional on being in the league last year: sample size: White (712), non White (444).

Table 2: Racial Differences in Outcomes:

	Coefficient	s.d.	R ²
Season fixed variable			
Ln initial Price	-0.019**	0.008	0.79
Never played during season	0.000	0.014	0.19
Duration of play during season	19.986	53.296	0.34
Total number of points during season	-2.118	2.240	0.37
Time varying variable			
Weekly change in Price	-0.0005	0.0005	0.092
Played this week	-0.003	0.018	0.122
Points accrue this week	-0.104	0.077	0.076
Minute played this week	-0.470	1.772	0.113

Note: Estimate of race effects on separate regressions.

The specification for season fixed variables includes: total points last year, value last year, not in league last year, season indicator, a quadratic in player's age, position, UK nationals, international status (English and other), club indicator. The sample size is then 1327 and the standard errors are clustered at the player level.

The specification for time varying variable include: player's characteristics, season, game week and number of games played in that week (0, 1, or 2). For weekly change in value, we also include a quadratic in points accrue in the previous week and initial value. Standard errors are then clustered at the player level and sample size is 58,759.

Table 3: OLS- Log Hires at the Start of the Season

	(1)	(2)	(3)	(4)	(5)
Non-White	0.025 (0.127)	0.095 (0.119)	0.214** (0.102)	0.073 (0.087)	0.102 (0.087)
Initial Price			1.396*** (0.181)		0.152 (0.193)
Initial Price Square			-0.050*** (0.010)		0.002 (0.011)
Points previous season				0.025*** (0.002)	0.025*** (0.003)
Points previous season ² / ₁₀₀₀				-0.024* (0.012)	-0.042*** (0.005)
Not in league last year				0.785*** (0.180)	0.738*** (0.185)
Constant	10.333*** (0.080)	-0.424 (1.747)	-0.119 (1.504)	5.429*** (1.313)	5.467*** (1.447)
Controls?	No	Yes	Yes	Yes	Yes
Observations	883	883	883	883	883
R-squared	0.00	0.21	0.37	0.50	0.51

Note: OLS estimates: Controls are: season dummy; player's position; quadratic in age; UK born; England international; international of another country, and club at which playing. Standard errors clustered at the individual level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Table 4: Quantile Regression of log Initial Hire

	Q10	Q25	Q50	Q75	Q90
Non-White	0.142 (0.148)	0.128 (0.127)	0.162 (0.105)	0.068 (0.130)	-0.256 (0.174)
Initial Price	0.617* (0.409)	0.653** (0.277)	0.190 (0.264)	-0.148 (0.274)	0.002 (0.313)
Initial Price Square	-0.023 (0.023)	-0.028* (0.017)	-0.003 (0.016)	0.018 (0.016)	0.007 (0.018)
Points previous season	0.020*** (0.004)	0.025*** (0.004)	0.029*** (0.005)	0.030*** (0.004)	0.029*** (0.004)
Points previous season ² / ₁₀₀₀	-0.009 (0.024)	-0.034* (0.019)	-0.056** (0.022)	-0.081*** (0.021)	-0.076*** (0.021)
Not in league last year	0.138 (0.238)	0.423 (0.263)	0.845*** (0.310)	0.959*** (0.293)	1.228*** (0.346)
Constant	2.517 (2.273)	5.044*** (1.804)	6.105*** (1.17)	7.000*** (2.107)	9.637*** (2.412)
Pseudo R-squared	0.33	0.35	0.35	0.31	0.27

Note: Quantile estimates are estimated simultaneously: Controls are: season dummy; player's position; quadratic in age; UK born; England international; international of another country, and club at which playing. Standard errors are obtained by bootstrap (100 replications)

*** p<0.01, ** p<0.05, * p<0.1. Sample size: N=883.

Table 5: OLS – Log Initial Hire – Test for Heterogeneity

Population	Non-White	R ²	N
Played in FF last season	0.119 (0.091)	0.57	706
Did not Play in FF last season	0.084 (0.307)	0.35	177
Points last season: Quartile 1	0.625*** (0.232)	0.41	137
Points last season: Quartile 2	0.209 (0.173)	0.36	190
Points last season: Quartile 3	0.217 (0.238)	0.35	180
Points last season: Quartile 4	0.014 (0.135)	0.44	199
Defender	0.039 (0.147)	0.48	304
Forward	0.037 (0.238)	0.63	154
Goalkeeper	-0.001 (0.529)	0.74	99
Midfielder	0.203 (0.147)	0.53	326
UK national	0.054 (0.148)	0.52	402
Non-UK national	0.120 (0.113)	0.54	481

Note: The non-White column reports the estimates on the coefficient for non-white players in a log initial hire model whose specification is identical to the one defined in Table 3, column(5). Standard errors, in parentheses, account for clustering at the player level: *** p<0.01, ** p<0.05, * p<0.1

Table 6: OLS - Net Employer Demand for Players

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	OLS	OLS	OLS	OLS	FE
Non White	-51.61 (93.53)	41.06 (103.25)	-71.68 (105.59)	-106.31 (137.51)	-1976.3* (1144)	
Points last game		1244*** (101.26)	1244*** (101.33)	1323*** (125.32)	1229*** (137.84)	1166*** (133.73)
-- *non-white					41.22 (194.40)	25.42 (186.66)
Points in previous 3 games		251.9*** (39.68)	265.7*** (39.67)	302.7*** (51.86)	249.8*** (52.13)	267.2*** (52.51)
-- *non-white					44.65 (75.09)	78.09 (78.39)
Total Points Accumulated		-60.84*** (5.49)	-49.39*** (5.25)	-40.49*** (7.71)	-46.78*** (6.41)	0.54 (8.28)
-- *non-white					-7.38 (9.02)	-22.82* (12.72)
Player price after last game			-4822.2*** (581.69)	-4585.8*** (705.79)	-5217.9*** (666.98)	-77665*** (7231.68)
-- *non-white					1130.48 (761.7)	-7109.70 (10363.1)
N picks for starting squads				-383.51*** (73.60)		
Constant	125.6* (69.75)	3962.9*** (1274.36)	8139.7*** (1395.21)	9953.4*** (1759.6)	8715.9*** (1472.21)	131455*** (10451)
Controls	No	Yes	Yes	Yes	Yes	Yes
R-squared	0.00	0.11	0.12	0.11	0.12	0.13

Note: Controls are season dummy; player's position; quadratic in age; UK born; England international; international of another country, club at which playing, game week, indicators of number of games played in the week. Nbr of observations = 58,319 or 38,511 in (4). Standard errors are clustered at the football player level. *** p<0.01, ** p<0.05, * p<0.1

Table 7: Quantile Regression: Net Employer Demand for Players

	Q10	Q25	Q50	Q75	Q90
Non-White	-19.05 (20.42)	2.72 (7.03)	-1.73 (2.40)	1.06 (1.39)	1.99** (0.88)
Ln Price	-10,703*** (371,70)	-2,743*** (115.06)	-455.7*** (28.25)	24.37** (10.79)	79.60*** (13.10)
Points	463.77*** (19.76)	223.18*** (7.27)	162.42*** (4.55)	649.76*** (24.45)	1,562*** (95.31)
Points in previous 3 games	-24.18* (13.12)	4.22 (3.78)	11.25*** (1.50)	52.43*** (5.08)	424.60*** (28.29)
Points accumulated up to t-4	-60.85*** (2.69)	-23.88*** (0.87)	-7.39*** (0.25)	-3.64*** (0.25)	-1.10*** (0.20)
Constant	14,749*** (598)	3,615*** (174)	702*** (46)	130*** (20)	60.57 (20)
Pseudo R-squared	0.17	0.06	0.01	0.08	0.27

Note: Quantile estimates are estimated simultaneously: Controls are season dummy; player's position; quadratic in age; UK born; England international; international of another country, club at which playing. game week and indicators of number of games played in the week. Nbr of observations = 58,319

Standard errors are obtained by bootstrap (100 replications)

*** p<0.01, ** p<0.05, * p<0.1.

Table 8: OLS - Net Employer Demand for Players – Test for Heterogeneity

Population	Non-White	R ²	N
Played at least once during season	-92.834 (114.22)	0.12	52,275
Not in league last year	-308.858** (153.68)	0.14	16,913
Initial Selection: Quartile 1	36.405 (33.30)	0.07	9,556
Initial Selection: Quartile 2	-101.789 (72.43)	0.07	9,654
Initial Selection: Quartile 3	-287.846** (120.12)	0.13	9,673
Initial Selection: Quartile 4	470.882 (474.40)	0.18	9,628
Defender	-99.478 (117.64)	0.14	19,425
Forward	-423.888 (273.14)	0.13	10,907
Goalkeeper	115.86 (229.45)	0.13	6,516
Midfielder	41.146 (194.35)	0.13	21,471
UK national	6.004 (131.18)	0.12	26,005
Non-UK national	-109.24 (146.076)	0.12	32,314
Week 2-10	263.03 (355.26)	0.23	12,898
Week 11-19	-300.75 (241.86)	0.12	13,746
Week 20-29	-315.18 (227.80)	0.10	16,202
Week 30-38	89.81 (198.18)	0.08	15,473

Note: The non-White column reports the estimates on the coefficient for non-white players in a net-transfer model whose specification is identical to the one defined in Table 5, column(3). Standard errors, in parentheses, account for clustering at the player level: *** p<0.01, ** p<0.05, * p<0.1

APPENDIX

TABLE A1: POINTS SYSTEM IN FANTASY FOOTBALL

ACTION	POINTS
Playing up to 60 minutes	1
Playing 60 minutes or more	2
For each goal scored by a goalkeeper or defender	6
For each goal scored by a midfielder	5
For each goal scored by a forward	4
For each goal assist	3
Clean sheet by a goalkeeper or defender	4
Clean sheet by a midfielder	1
For every 3 shot saves by a goalkeeper	1
For each penalty save	5
For each penalty miss -	-2
For every 2 goals conceded by a goalkeeper or defender	-1
For each yellow card	-1
For each red card	-3
For each own goal	-2
Bonus points for the best players in a match	1/3

BONUS POINTS:

The three best performing players in each match according to the Bonus Points System will receive additional bonus points. 3 points will be awarded to the highest scoring player, 2 to the second best and 1 to the third. The Bonus Points System is based on statistics on various dimensions of the player's performance collected by OPTA (<http://www.optasports.com/>).

APPENDIX FIGURE A1

The screenshot shows the Fantasy Premier League website interface. At the top, there are navigation tabs for 'PITCH VIEW' and 'DATA VIEW'. Below this, the user's name 'RnosDream' is displayed. The main content area features a detailed profile for Emmanuel Adebayor, valued at £9.2m. This profile includes a 'Player summary' with his position (Forward) and team (Tottenham), a 'Recent matches' table, and a 'Gameweek 31' performance table. To the right, there are 'Player filters' and an advertisement for an iRobot vacuum. At the bottom, a 'Gameweek 32 Fixtures' section shows a match between Swansea and Newcastle. A sidebar on the right lists top-performing players like Lampard, Sinclair, and Valencia.

Emmanuel Adebayor £9.2m

Player summary

Position: Forward
 Team: Tottenham
 Teams selected by: 10.9%
 Total score: 160

Recent matches

GW	Opponent	Points
29	STO (H)	0
30	CHE (A)	2
31	SWA (H)	13

Gameweek 31

Statistic	Value	Points
Bonus	3	3
Minutes played	90	2
Goals scored	2	8
Total		13

Upcoming fixtures

GW	Opponent	Date
32	SUN (A)	07 Apr
33	NOR (H)	09 Apr
34	-	-

Gameweek 32 Fixtures

06 Apr 16:30 Swansea v Newcastle

Top Performers

Player	Team	Points
Lampard	CHE	11.0
Sinclair	SWA	6.7
Valencia	MUN	8.2
Larsson	SUN	6.7
Walters	STO	6.3
Arteta	ARS	7.9
Forwards		6.73
Van Persie	ARS	13.6
Rooney	MUN	13.0
Adebayor	TOT	9.2
Aguero	MCI	10.8

1